

Final Cooling with Thick Wedges for a Muon Collider

Daniel Fu^{1, a)}

University of Chicago

(Dated: 11 August 2023)

A muon collider is a particle accelerator that collides muons rather than protons or electrons. An important component of such a collider is cooling, in which the emittance of the muon beam is reduced to make it suitable for use in a collider. In the last stage of this process, called final 4D cooling, emittance in the transverse axes is reduced while emittance in the longitudinal axis is allowed to grow. Final 4D cooling is poorly understood, and previously published cooling channels fail to achieve the required emittances without serious technical challenges, such as the requirement for high-field solenoids. In this study, we developed and optimized a conceptual design for the 4D cooling channel based on single thick wedges. We used G4Beamline to simulate the channel and Python to generate and analyze particle distributions. We characterized the effects of various parameters on the performance of this system, including the initial emittances and Twiss parameters, the length and angle of the wedges, the length of the drift channel, and the length, frequency, gradient, and phase of the RF cavity. We optimized the design parameters of the cooling channel and produced two conceptual designs (corresponding to two possible starting points for the input beam) which achieve transverse cooling by a factor of 3.5. These channels achieve a lower transverse and longitudinal emittance than the best design previously published.

^{a)}danielfu@uchicago.edu

I. INTRODUCTION

A muon collider is a type of particle accelerator that collides muons, rather than protons or electrons. Muons have several properties that make them desirable for accelerator research. Like electrons, they are elementary particles, making collisions between them simpler to model and more efficient. Like protons, they are more massive than electrons, reducing losses due to synchrotron radiation and allowing much higher energies to be attained. These properties would allow a relatively small muon collider (around 2 km diameter) to have physics reach comparable to much larger proposed colliders (such as the 30 km ILC and 50 km CILC).

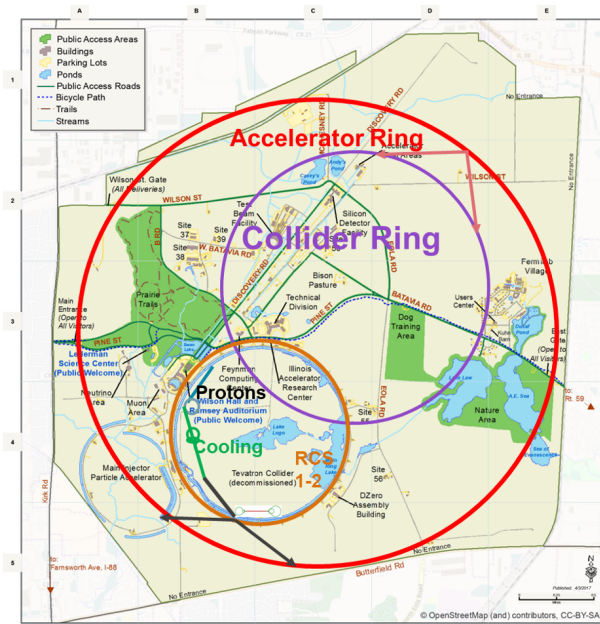


FIG. 1. Diagram of how a muon collider could be built on the Fermilab site

The main disadvantage of muons is that they are unstable. Thus, unlike protons and electrons, which can be collected from the environment, muons are produced in collisions of a high-energy proton beam with a stationary target. These collisions produce pions, which in turn decay into muons. These muons are produced with a high emittance (spread of positions and momentums) which must be reduced in a process known as cooling before they can be used in an accelerator. In the final stage of this cooling process, the emittance of the beam in the transverse axes (those perpendicular to the direction of travel) must be reduced to below $30 \mu\text{m}$, while the emittance in the longitudinal (parallel to direction of travel) axis is allowed to increase; this process is termed 4D cooling, as it reduces spread in two axes of position and two axes of momentum.

In this paper, we conduct a simulation-based investigation of a final 4D cooling method proposed by David Neuffer¹. This cooling method passes the beam through a wedge of solid material, which absorbs energy from the muons through ionization. Because different parts of the beam pass through different amounts of material, this wedge introduces dispersion, correlating the position of particles along one transverse axis with the momentum of those particles. This dispersion is




FIG. 2. Path of the muon cooling process in transverse and longitudinal emittances. The final 4D cooling step we investigated is circled in blue. The previous best result is marked in red, and our best result is marked in blue.

then removed using a bending magnet, which displaces the particles towards the centerline of the beam, reducing the emittance in that axis. This process increases the spread in momentums of the muons (measured by the standard deviation of momentum, σ_p), so the beam is then sent through phase rotation to decrease σ_p at the cost of increasing σ_t (the spread of the particles in time). The beam is then passed through a second wedge to reduce the emittance of the other transverse axis. The layout of the cooling channel is shown in Fig. 3.

We characterized the effects of various parameters of the input beam, wedges, and phase rotation on the performance of the final cooling system and produced an optimized design for the channel. We considered two starting points for our final cooling channel: 145 μm transverse emittance with 1.0 MeV/c σ_p , and 110 μm transverse emittance with 0.8 MeV/c σ_p . These correspond




FIG. 3. Diagram of the wedge cooling pathway investigated.

to the outputs of the B10 and B11 stages in a previously designed ionization cooling channel using high-field solenoids². Our channel could therefore replace the final stages of this design, further reducing the transverse emittance while avoiding the need for extremely high-field solenoids.

II. METHODS

We wrote a Python script to generate particle distributions with specified mean momentum (p), emittances ($\epsilon_x, \epsilon_y, \epsilon_z$), Twiss parameters (β, α, γ) in the transverse axes, and standard deviation of momentum (σ_p), following an algorithm described by Mike Syphers³.

We used G4Beamline, a Geant4-based simulation software, to model the effects of the wedges and RF cavity on these distributions. The geometry of the wedges is defined by two parameters: the wedge length (measured at the point where the centerline of the beam crosses the wedge) and the half-angle. Fig. 4 shows how these parameters are defined. Diamond was used as the wedge material for all runs.

We used a virtual detector (VD) to get results from the simulation. This is a plane defined in G4Beamline where the positions and momentums of the particles are measured. The center of the wedge is positioned halfway between the beam source and the virtual detector. The distance to the virtual detector was held constant across runs. As Fig. 4 shows, we also set the base length of the wedge to this distance. An excessively small VD distance would lead to the wedge being cut off early and particles bypassing the wedge. As demonstrated in Fig. 5, any VD distance greater than about 3 times the wedge length produces roughly the same result. As we were working with




FIG. 4. Diagram of the setup used in simulations, showing wedge (purple), muon beam (orange), and virtual detector (light blue), with wedge length, wedge half-angle, and virtual detector distance marked. Note that wedge length is defined as the distance that the beam centerline intersects the wedge.




FIG. 5. Graph of ϵ_x v.s. distance to the virtual detector. This graph was made with a wedge of length 7.5 mm, and the red dotted line marks a VD distance of three times this length. Note that the curve levels off before this point. Error bars of 1 standard deviation are marked.

wedge lengths less than 8 mm, we set VD distance to a constant 24 mm across all runs.

We calculated the transverse emittances and Twiss parameters from G4Beamline output files using the following formulas:

$$\begin{aligned}\epsilon\beta &= \langle x^2 \rangle - \frac{\langle x\delta \rangle^2}{\langle \delta^2 \rangle} \\ \epsilon\gamma &= \langle x'^2 \rangle - \frac{\langle x'\delta \rangle^2}{\langle \delta^2 \rangle} \\ \epsilon\alpha &= -\langle xx' \rangle + \frac{\langle x\delta \rangle \langle x'\delta \rangle}{\langle \delta^2 \rangle} \\ \epsilon &= \sqrt{(\epsilon\beta)(\epsilon\gamma) - (\epsilon\alpha)^2}\end{aligned}$$

Dispersion was calculated using the following formulas:

$$\begin{aligned}D &= \frac{\langle x\delta \rangle}{\langle \delta^2 \rangle} \\ D' &= \frac{\langle x'\delta \rangle}{\langle \delta^2 \rangle}\end{aligned}$$

Designing the dispersion correction system and transport/focusing lattice was beyond the scope of this project, so we assumed idealized versions of these components. We model the removal of




FIG. 6. Flowchart showing our simulation process. Steps simulated in G4Beamline are highlighted in green, steps mathematically modeled are highlighted in orange.

dispersion as a linear subtraction of the calculated dispersion values:

$$x_0 = x - D\delta$$

$$x'_0 = x' - D'\delta$$

We simulated only longitudinal momentum through the phase rotation (drift channel and RF cavity), making the assumption that the beam can be transported through these components without significantly distorting the transverse momentum. The drift channel was modeled by adding the time each particle would take to transverse it (based on the particle's p_z) to that particle's time coordinate. The p_x and p_y of particles were set to zero before they were passed through the RF cavity. The x and y momenta were then re-added before the distribution was passed into the second wedge.

The beam exiting the first wedge is diverging, as the initial distribution is focused to a point within the first wedge. To model refocusing of the beam prior to entering the second wedge, we inverted the x and y momenta of all particles. This had the effect of negating α . This represents a rough approximation of refocusing between the first and second wedges. A real focusing system would be able to set β and α to optimal values before the second wedge, and could therefore achieve better performance.

We began by identifying variables that were monotonically correlated with wedge performance and thus were constrained by technical limitations of the cooling system's construction. These were the initial momentum, σ_p , Twiss parameter β , and transverse emittance, the length of the drift channel, and the frequency of the RF cavity. We then optimized the remaining parameters using the the Nelder-Mead method as implemented in Scipy.




FIG. 7. ϵ_x after first wedge v.s. wedge length and half-angle for the first and second wedges in the $145 \mu\text{m}$ case, with optimal values marked.

G4Beamline runs for the optimization and graphs were run with 12,000 particles. More particles (200,000) were used for the phase-space plots and statistics of the optimal channel.

III. RESULTS

A. Wedges

We studied the effects of wedge geometry (centerline length and half-angle), Twiss parameters, momentum, σ_p (standard deviation of momentum), and initial emittances.

In all these plots, we applied a 4σ cut to the output distribution (removing outliers that were more than 4 standard deviations from the mean in x , y , t , p_x , p_y , and p_z).

1. Wedge geometry

We found that there exists an optimal combination of wedge length and wedge half-angle, which varies depending on the momentum of the particles and Twiss parameters of the distribution, among other factors. Fig. 7 demonstrates that the optimal geometries of the first and second wedges are different due to the different particle distributions entering the two wedges.

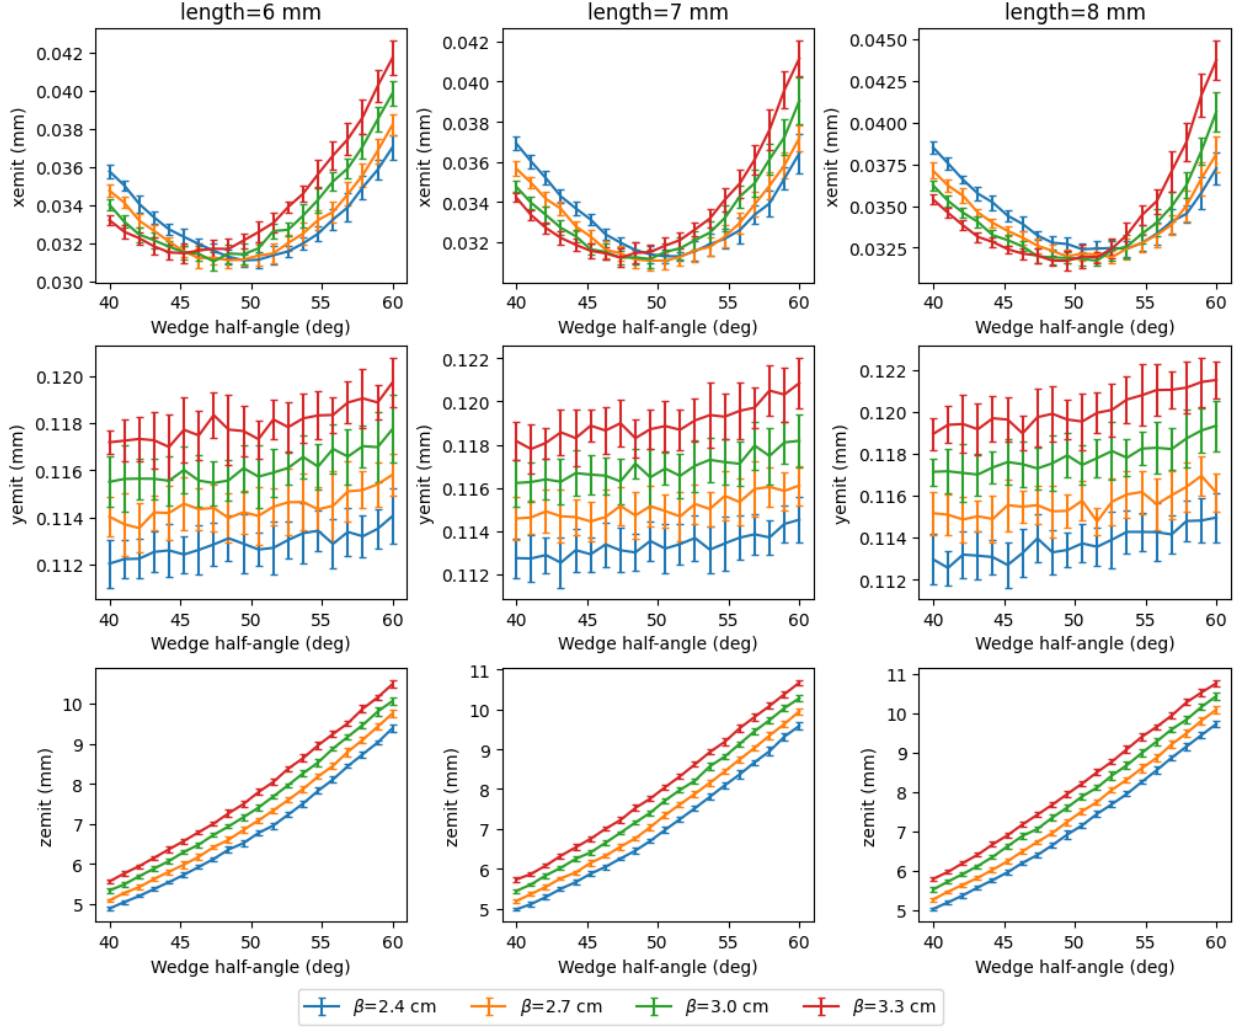


FIG. 8. Emittances after first wedge v.s. half-angle for various β and wedge length in the $110\ \mu\text{m}$ case; error bars of 1 standard deviation are shown. Note that $\alpha = 0.7$ for this graph, while we used $\alpha = 1.0$ in the final optimization of this case.

2. Twiss parameters

As shown in Fig. 8, increasing β decreases the optimal half-angle and increases ε_y (that is, the transverse emittance in the axis the wedge is not acting in). Very high or low values of β increase ε_x , but there is a significant range of β that produces about the same minimum ε_x value. This range appears to move towards larger β with increasing length. Table I gives the minimums of ε_x corresponding to four values of β .

As Fig. 9 shows, lower values of β also decrease output σ_p . As both increased ε_y and σ_p going into the second wedge will increase ε_y after the second wedge, β should be made as low

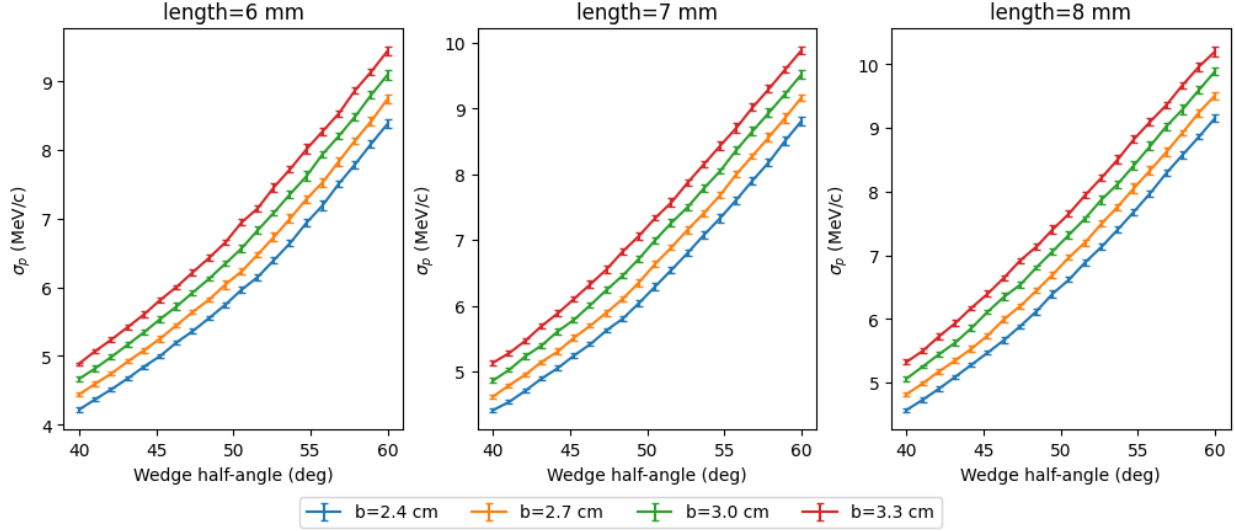


FIG. 9. Post wedge σ_p v.s. angle for various β and wedge length in the $110\ \mu\text{m}$ case; error bars of 1 standard deviation are shown. $\alpha = 0.7$ for this graph as well.

| β (cm) | ϵ_x (μm) | ϵ_y (μm) | ϵ_z (mm) | angle (deg) | length (m) |
|--------------|--------------------------------|--------------------------------|-------------------|-------------|------------|
| 2.4 | 31.1 | 112.9 | 6.52 | 49.5 | 6 |
| 2.7 | 31.1 | 114.9 | 7.35 | 50.5 | 7 |
| 3.0 | 31.0 | 115.5 | 6.73 | 47.4 | 6 |
| 3.3 | 31.2 | 119.0 | 7.22 | 47.4 | 7 |

TABLE I. Emittances after the first wedge and wedge parameters at minimum ϵ_x for various values of β in the $110\ \mu\text{m}$ case.

as possible without significantly reducing first-wedge performance. Lower β corresponds to more powerful focusing magnets, so a minimum β is primarily set by technical limitations. We chose $\beta = 3.0\text{cm}$ as a reasonable value that is within the optimal range for the first wedge.

As shown in Fig. 10, increasing α decreases ϵ_y . There is also a value for α that minimizes ϵ_x , above or below which ϵ_x is increased. This optimal value becomes lower as the wedge length is increased. Increasing α is thus a trade-off between minimizing ϵ_x and ϵ_y . Table II gives the minimums of ϵ_x corresponding to four values of α .

We chose $\alpha = 0.7$ for the $145\ \mu\text{m}$ case, as it is close to optimal for the first wedge's length. The $110\ \mu\text{m}$ case had problems with higher final ϵ_y than desired, so we switched to $\alpha = 1.0$ for this

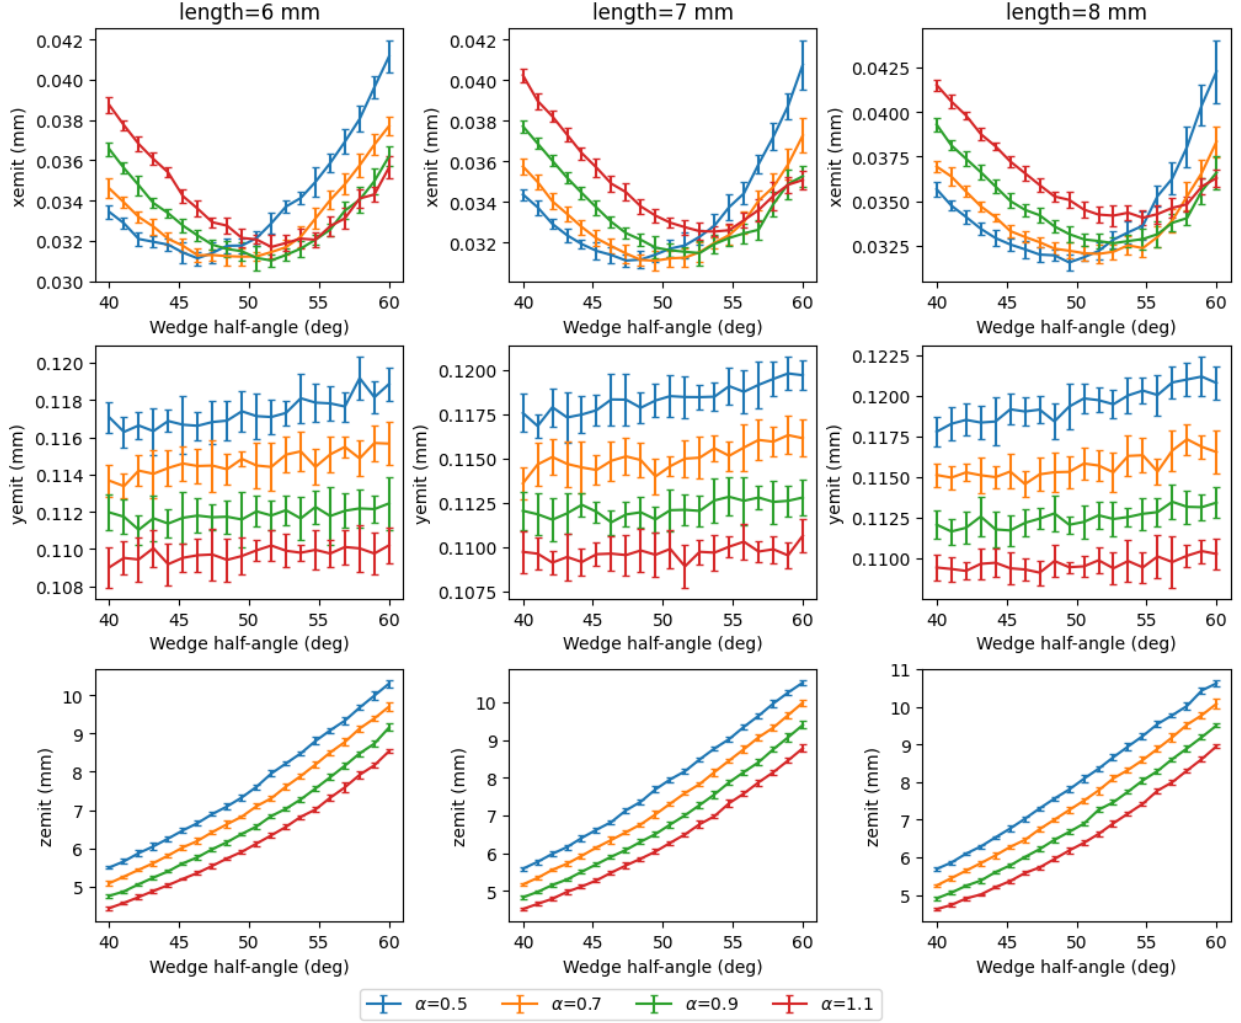


FIG. 10. Emittances after first wedge v.s. half-angle for various α and wedge length in the $110\ \mu\text{m}$ case; error bars of 1 standard deviation are shown. Note that $\beta = 2.7\ \text{cm}$ for this graph, while we used $\beta = 3.0\ \text{cm}$ in the final optimization of this case.

| α | ε_x (μm) | ε_y (μm) | ε_z (mm) | angle (deg) | length (m) |
|----------|-----------------------------------|-----------------------------------|----------------------|-------------|------------|
| 0.5 | 31.1 | 118.3 | 7.14 | 47.4 | 7 |
| 0.7 | 31.1 | 114.0 | 7.03 | 49.5 | 7 |
| 0.9 | 31.0 | 111.8 | 6.84 | 51.6 | 6 |
| 1.1 | 31.7 | 110.2 | 6.34 | 51.6 | 6 |

TABLE II. Emittances after the first wedge and wedge parameters at minimum ε_x for various values of α in the $110\ \mu\text{m}$ case.




FIG. 11. Minimum ϵ_x v.s. σ_p for the 110 μm case, as well as optimal length and half-angle. Note that length and half-angle were re-optimized for each trial. Error bars of 1 standard deviation are shown.

case.

3. Momentum spread

Fig. 11 shows that as σ_p is increased, the minimum achievable ϵ_x increases (assuming wedge geometry is re-optimized for each σ_p). Minimizing σ_p is thus important to improve the performance of the wedge. Increasing σ_p also increases the optimal half-angle, but seems to have no significant effect on the optimal length. We assumed a σ_p of 1.0 MeV/c for the 145 μm case and 0.8 MeV/c for the 110 μm case.

4. Initial transverse emittance

As shown in Fig. 12, lower initial transverse emittance results in lower post-wedge ϵ_x , as expected. Increasing initial transverse emittance also decreases the optimal half-angle and slightly increases the optimal wedge length.

5. Momentum

As shown in Fig. 13, higher input momentum increases ϵ_x . This correlation appears to be levelling off around 100 MeV/c in the 7 mm case, indicating that there might be a minimum




FIG. 12. Minimum ϵ_x v.s. initial transverse emittance as well as optimal length and half-angle. $\sigma_p = 0.8$ MeV/c, $\beta = 3.0$ cm, and $\alpha = 0.7$ were kept constant. Length and half-angle were re-optimized for each trial. Error bars of 1 standard deviation are shown.

of ϵ_x that moves towards higher momentum with higher wedge length. Assuming it exists, this minimum is below the range of momenta considered achievable from the previous stage of cooling (around 100-120 MeV/c). Lower momentum also increases ϵ_z and has no significant effect on ϵ_y . Given that ϵ_z can be allowed to grow significantly in 4D cooling, we conclude that input momentum should be as low as possible. We used a momentum of 100 MeV/c for all runs.

B. Phase rotation

1. Drift length

Fig. 14 demonstrates that a longer drift length produces better phase rotation performance (lower σ_p). As shown in Fig. 15, a longer drift length produces a greater spread in time for phase-rotation to act upon. The constraint on drift length is thus how long of a channel can be feasibly constructed. We chose 16 m as a reasonable value and used it for all runs. This length was measured from the start of drift to the center of the RF cavity.

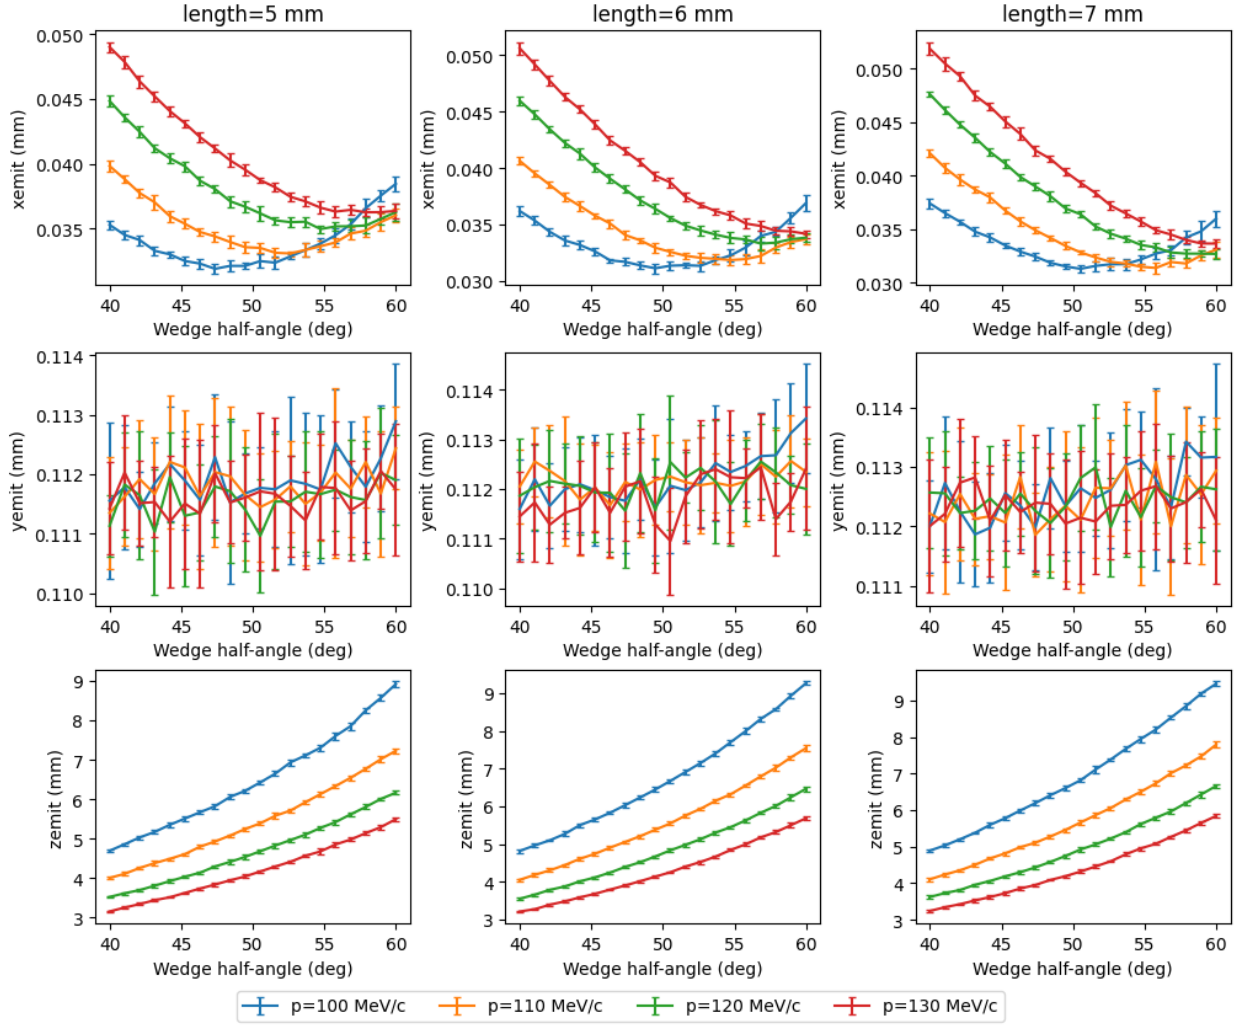


FIG. 13. Emittances after first wedge v.s. half-angle for various beam momentums and wedge length in the 110 μm case; error bars of 1 standard deviation are shown.

2. RF phase

Nelder-Mead optimization gave a value of 0 degrees as the optimal phase angle of the RF field. As show in 16, the optimal phase angle is actually slightly below 0. This difference was presumably too small for optimization to pick up on, so while we've used 0 as the phase angle of the field in all future simulations, the performance of the second wedge could be somewhat improved by further optimizing the RF phase.




FIG. 14. Best achievable σ_p , and optimal values of RF cavity length and gradient, in the 145 μm case. RF frequency and phase are set to constant values as described further below, while RF cavity length and gradient are re-optimized at each point. No error bars are shown, as only one trial at each σ_p was performed.




FIG. 15. Longitudinal phase-space distribution after various lengths of drift in the 145 μm case.




FIG. 16. σ_p after phase rotation v.s. RF phase angle in the 145 μm case. Drift length, RF frequency, RF cavity length, and RF gradient are kept constant. Error bars of 1 standard deviation are shown.

3. RF frequency, cavity length and gradient

We found that the effects of RF frequency, RF cavity length, and RF gradient are dependent on each other. As Fig. 17 shows, the optimal gradient value falls and then rises with increasing frequency. Within this curved valley, the optimal frequency is at around 20 MHz. However, constructing a RF cavity with frequency this low was considered impractical. We selected 25 MHz as the lowest frequency deemed achievable and used this frequency in all runs.

Fig. 18 shows that RF gradient has a similar relationship with RF cavity length. The optimal values of these two parameters (assuming a frequency of 25 MHz) were found using Nelder-Mead optimization.

Fig. 19 shows that a higher frequency causes distortions in the phase space. These are the result of the actual RF oscillation. A lower frequency is thus desired so that most of the particles remain in the area of the RF oscillation where the field is linear with time.

Fig. 20 shows that the RF gradient changes the intensity of the cavity's effect on the distribution, effectively rotating the phase-space more or less. It should thus be tuned to rotate the phase-space to horizontal (minimal σ_p).

Fig. 21 shows that the length of the cavity also affects the amount of rotation, as well as adding

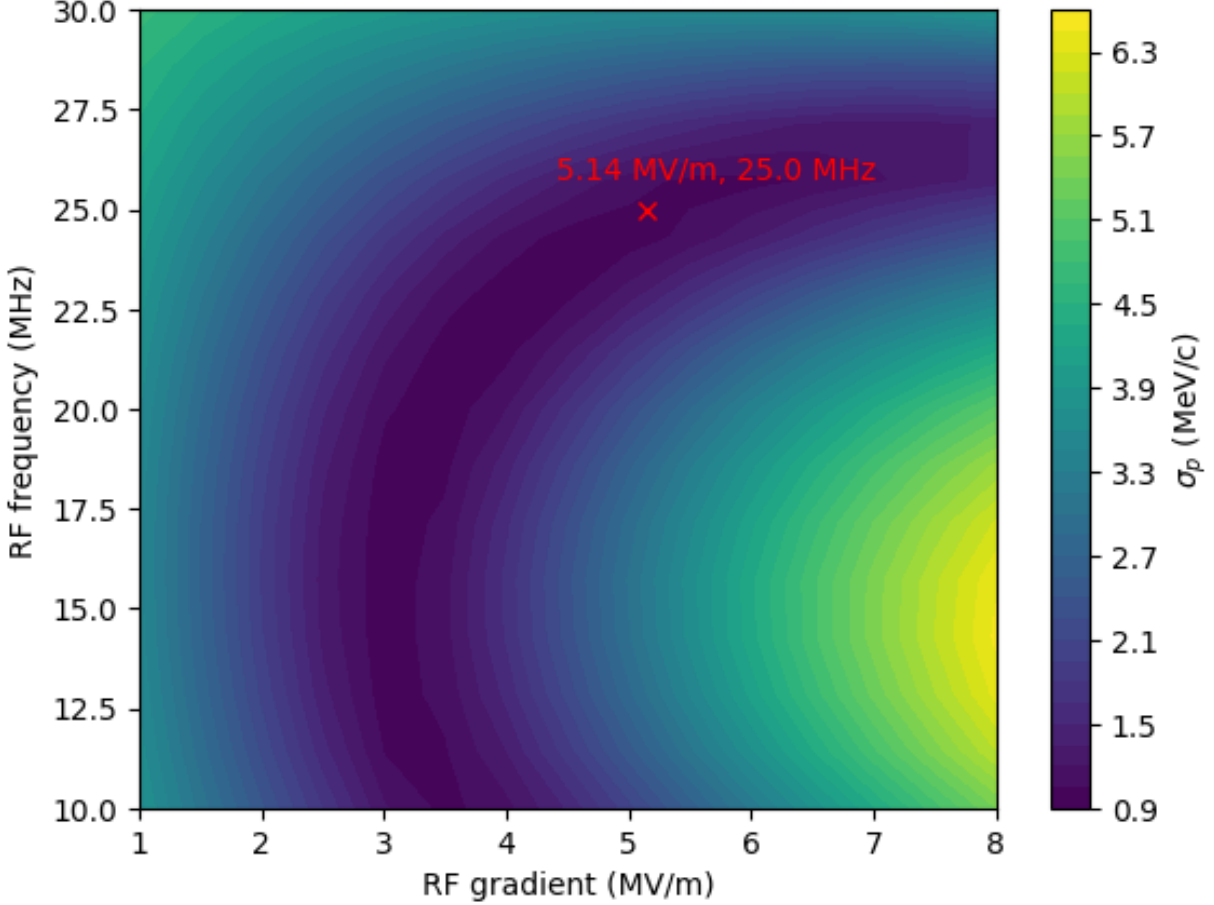


FIG. 17. Plot of final σ_p for different values of RF frequency and gradient in the 145 μm case, with the selected parameter pair marked.

some distortion to the phase-space. The similarity of this effect to that of gradient explains why they are linked. It can be seen that lengths in the 5 m-6 m range produce the least distortion.

Fig. 22 shows that RF phases off 0 cause a curvature in the phase-space distribution. This explains why the optimal phase angle is close to 0. Phase angles off 0 also accelerate or decelerate the particle bunch as a whole.

C. Full optimized channel

We produced an optimized channel for both the 145 μm and 110 μm cases. Table III gives the optimized parameters. Fig. 23 and Fig. 24 show the evolution of the phase-space distribution as the particles travel through the channels. Tables IV and V give the emittances, standard deviation of




FIG. 18. Plot of final σ_p for different values of RF cavity length and gradient in the 145 μm case, with the optimal parameters marked. Frequency is fixed at 25 MHz for this plot.

| Case | 145 μm | 110 μm |
|----------------------|-------------------|-------------------|
| Wedge 1 length (mm) | 7.5 | 6.6 |
| Wedge 1 angle (deg) | 46.5 | 51.6 |
| RF cavity length (m) | 5.6 | 5.7 |
| RF gradient (MV/m) | 5.1 | 5.1 |
| Wedge 2 length (mm) | 6.1 | 5.1 |
| Wedge 2 angle (deg) | 40.8 | 47.2 |

TABLE III. Optimal wedge lengths, wedge angles, RF lengths, and RF gradients for both cases considered.




FIG. 19. Longitudinal phase-space distribution after phase rotation for various RF frequencies in the 145 μm case. RF length and gradient are held constant.

momentum, standard deviation of time, average momentum, and beam remaining (transmission) at various points along these channels.

The first wedge in both cases achieves cooling by a factor of about 3.5 (145 μm to 40 μm and 110 μm to 32 μm). The second wedge performs less well than the first, only able to reduce ϵ_y to 43.5 μm in the 145 μm case and 33.8 μm in the 110 μm case. This is likely due to the higher σ_p going into the second wedge, as the phase-rotation only reduces σ_p to around 1.3 MeV/c (rather than the 1.0 MeV/c or 0.8 MeV/c that the beam began with). Improvements to the phase-rotation setup could therefore result in improved y-axis performance.

Two cuts are performed in this simulation: a 15% cut of momentum after the phase rotation, and a 4σ cut after the second wedge. The 4σ cut removes less than 0.5% of the beam, but decreases the emittances by a substantial amount, indicating that a significant part of the emittance was due to outlier particles far from the center of the distribution. The channels overall have a transmission of about 83.5%.




FIG. 20. Longitudinal phase-space distribution after phase rotation for various RF gradients in the 145 μm case. RF length and frequency are held constant.

| Stage | ϵ_x (μm) | ϵ_y (μm) | ϵ_z (mm) | σ_p (MeV/c) | σ_t (ns) | p (MeV/c) | Transmission |
|------------------------------|-----------------------------------|-----------------------------------|----------------------|-----------------------|--------------------|----------------|--------------|
| Initial distribution | 144.9 | 145.1 | 1.262 | 1.001 | 0.745 | 100.0 | 100.0% |
| After first wedge | 45.9 | 151.4 | 6.319 | 7.179 | 0.745 | 87.2 | 100.0% |
| After RF cavity + 15% cut | 37.7 | 139.5 | 4.401 | 1.304 | 4.041 | 86.3 | 84.0% |
| After second wedge | 44.5 | 50.1 | 29.772 | 8.063 | 4.040 | 72.6 | 83.8% |
| After 4σ cut | 40.1 | 43.9 | 28.649 | 7.787 | 4.038 | 72.7 | 83.4% |

TABLE IV. Beam parameters through the optimal channel for the 145 μm case. This simulation was run with 200,000 particles.




FIG. 21. Longitudinal phase-space distribution after phase rotation for various RF cavity lengths in the $145\ \mu\text{m}$ case. RF frequency and gradient are held constant.

| Stage | ϵ_x (μm) | ϵ_y (μm) | ϵ_z (mm) | σ_p (MeV/c) | σ_t (ns) | p (MeV/c) | Transmission |
|------------------------------|-----------------------------------|-----------------------------------|----------------------|-----------------------|--------------------|----------------|--------------|
| Initial distribution | 109.8 | 109.4 | 1.390 | 0.800 | 0.930 | 100.0 | 100.0% |
| After first wedge | 34.5 | 114.1 | 7.115 | 6.301 | 0.930 | 88.9 | 100.0% |
| After RF cavity + 15% cut | 27.9 | 105.4 | 4.886 | 1.434 | 3.611 | 87.9 | 84.0% |
| After second wedge | 32.4 | 39.4 | 24.542 | 6.838 | 3.611 | 77.0 | 83.9% |
| After 4σ cut | 29.9 | 36.3 | 23.646 | 6.632 | 3.608 | 77.1 | 83.6% |

TABLE V. Beam parameters through the optimized channel for the $110\ \mu\text{m}$ case. This simulation was run with 200,000 particles.




FIG. 22. Longitudinal phase-space distribution after phase rotation for various RF phases in the $145\ \mu\text{m}$ case. RF length, frequency, and gradient are held constant.

IV. CONCLUSION

Muon beam cooling is one of the critical problems in the design of a muon collider. The earlier stages of 6D cooling are relatively well understood, but the design of the final 4D cooling channel is less certain. The best previously published channel⁴ achieves $55\ \mu\text{m}$ transverse emittance with $76\ \text{mm}$ longitudinal emittance. In this paper, we have designed a channel that outperforms the Sayed channel in both transverse and longitudinal emittance, with a final transverse emittance below $35\ \mu\text{m}$ and longitudinal emittance of $24\ \text{mm}$. With the assumption that the dispersion correction and focusing lattice can be constructed to a high degree of performance, these channels represent a plausible design for final 4D cooling and a step towards the construction of a muon collider.




FIG. 23. Phase-space distributions at various points along the optimized channel for the 145 μm case.




FIG. 24. Phase-space distributions at various points along the optimized channel for the 110 μm case.

REFERENCES

- ¹D. Neuffer, H. Sayed, J. Acosta, T. Hart, and D. Summers, “Final cooling for a high-energy high-luminosity lepton collider,” JINST (2017).
- ²D. Summers, “More muon cooling, higher luminosity,” https://indico.cern.ch/event/961803/contributions/4064621/attachments/2141709/3608895/DS_mu_12Nov.pdf (2020).
- ³M. Syphers, “Creation and analysis of beam distributions,” <https://nicadd.niu.edu/~syphers/tutorials/analyzeTrack.html> (2019), last accessed 9 August 2023.
- ⁴H. K. Sayed, R. B. Palmer, and D. Neuffer, “High field – low energy muon ionization cooling channel,” PHYSICAL REVIEW SPECIAL TOPICS (2015).